WHITE PAPER

THE ROLE OF HIGHER BAUD RATES IN EVOLVING COHERENT TRANSPORT

Identifying the Benefits and Use Cases for Higher Baud Rates

Since its emergence in the late 2000s, coherent technology has undergone a number of transitions. First from 40G to 100G, then from hard decision FEC to soft decision FEC, and then from single rate interfaces to flexi-rate interfaces with additional modulation schemes for 150G (8QAM) and 200G (16QAM). With no end in sight to the traffic growth from video, cloud, and DCI, and with the impact of 5G yet to be known, network operators need to scale capacity cost effectively while minimizing power consumption and footprint. To address this need, network equipment vendors are taking the next steps to evolve coherent optics technology with the option of higher baud rates.

MARKET DRIVERS FOR COHERENT TRANSPORT EVOLUTION

IP traffic, both between data centers and users and directly between users, and data center interconnect traffic represent the vast majority of traffic running over optical networks today and in the future, as shown in Figure 1. This traffic is growing with a CAGR of over 25%, doubling every three years. In particular, DCI traffic is forecast to grow with a CAGR of 31.9%, internet video with a CAGR of 31%, and mobile data with a CAGR of 46%.

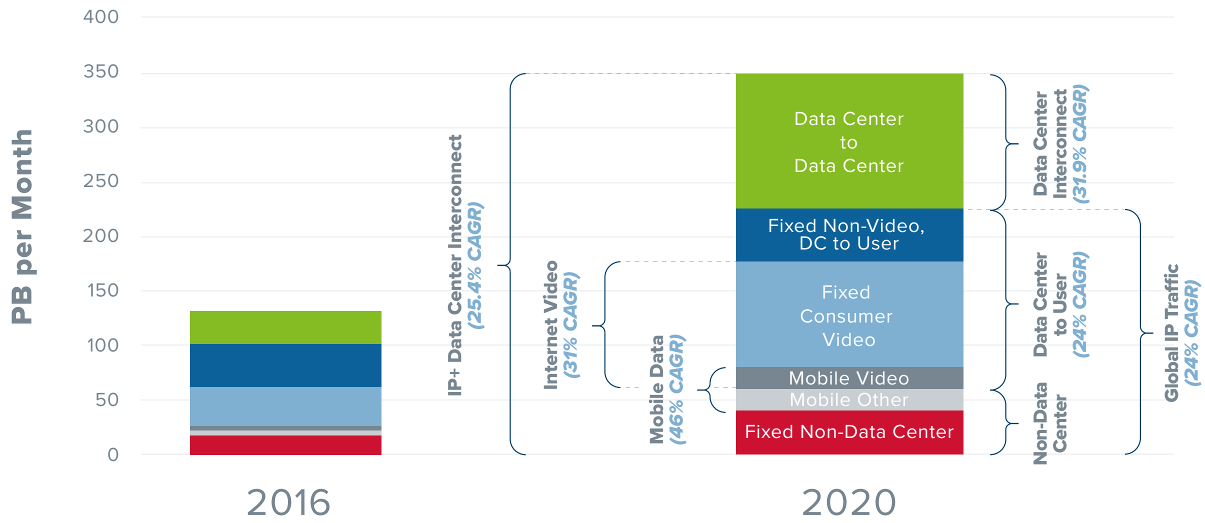

FIGURE 1 – *Optical Network Traffic: IP + DCI (Source: Cisco VNI 2017 & Cisco GCI 2016)*

Furthermore, it is important to note that these growth rates represent averages and may in fact be significantly underestimating the need for increased optical transmission capacity. For example, while average IP traffic is growing with a CAGR of 24%, busy hour traffic is growing at closer to 36%. This busy hour traffic growth has a far greater impact on the optical transport network, which is typically dimensioned for peaks rather than average traffic loads. Evolving coherent technology has a key role to play if this bandwidth escalation is to be delivered without dramatic increases in both CapEx and OpEx.

CapEx: Cost Per Bit and Spectral Efficiency

The cost per bit is the most direct driver of optical network CapEx. The key to lowering the cost per bit is achieving more bits per second with the same hardware investment in DSPs and optical transceiver components. In long haul applications, reach also becomes a key driver of the cost per bit enabling network operators to avoid costly OEO regens and to operate at higher wavelength capacities. Spectral efficiency is another driver of network CapEx. As bandwidth increases and strains the capacity of both fiber and the WDM line systems, increased spectral efficiency can avoid, or at least delay, the CapEx investment required to upgrade the WDM line system and/or light new fiber.

OpEx: Power Consumption and Footprint

Power consumption and footprint are key drivers of network operational costs. Today's state-of-the-art optical transmission can deliver 0.2 W per Gbps and 3.2 Tbps (1.6T line + 1.6T client) in 1RU. Further decreasing power consumption and footprint are key to containing operational costs while meeting the need for a doubling of transport capacity every three years.

Open Line Systems and Disaggregation

A key industry trend in optical transport is disaggregating DWDM systems into Open Line Systems (OLSs) and traffic bearing functional blocks such as high-capacity muxponders (refer to the Infinera white paper *The Case for Open Line Systems*). This trend promises reduced vendor lock-in, competitive pricing, and faster innovation leading to both lower CapEx and lower OpEx. Disaggregating DWDM systems into OLSs creates a requirement for the next generation of coherent transport to provide options that can deliver optimal performance for a given third-party OLS, whatever its capabilities. A secondary consideration for disaggregated DWDM networks is the ability to interwork the coherent optics of different vendors leveraging common modulation and Forward Error Correction (FEC).

New Client Interface Types: 200 GbE, 400 GbE, FlexE, and FlexO

Higher speed client interfaces enable network operators to lower the cost per bit, increase router faceplate density, and overcome the limitations of today's schemes for balancing traffic over a number of interfaces such as 802.1AX Ethernet link aggregation and IP Equal Cost Multi-Path (ECMP). This coupled with the standardization and availability of pluggables, including QSFP56 for 200 GbE and CFP8, QSFP-DD, and OSFP for 400 GbE, is expected to drive demand for 200 GbE from 2019 and 400 GbE from 2021, as shown in Figure 2.

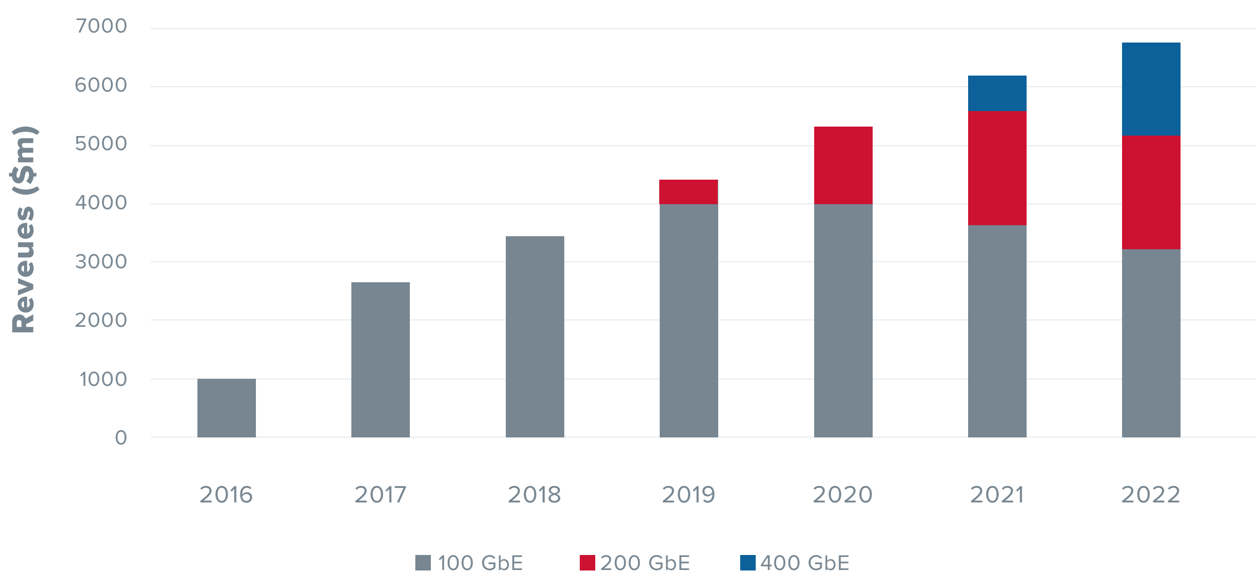

FIGURE 2 – *100 GbE, 200 GbE, and 400 GbE Transceiver Revenues (Ovum, May 2017)*

Additional new client standards include Flexible Ethernet (FlexE) and Flexible OTN (FlexO). With an interoperability agreement published by the OIF in 2016, FlexE enables more efficient bonding over multiple links and the sub-rating of high speed links, all with granularity down to as little as 5 Gbps. Specified in G.709.1 in early 2017, FlexO provides 100G granularity for short reach OTN clients, reusing multiple existing 100G pluggables (CFP2, CFP4, QSFP28, etc.). The need to efficiently transport these new clients on a single lambda provides an additional driver for evolving coherent transport.

HIGHER BAUD RATES VS. HIGHER ORDER MODULATION

These market drivers are leading optical equipment vendors to evolve coherent transport to support higher capacity wavelengths. They have two primary levers for doing this: the baud rate, the number of symbols per second, and the modulation, the number of bits per symbol. Each has its advantages, as shown in Table 1.

		Higher Order Modulation	Higher Baud Rate
Increased Spectral Efficiency		✓	**x** *(with exceptions)*
Reach Reduction with 2x Capacity		~75%	~10%
Lower Cost Per Bit		✓	✓
Applicability	Flexi-grid Networks	✓	✓
	Fixed Grid (50 GHz) Point-to-point	✓	✓ *(limited max baud rate)*
	Fixed Grid (50 GHz) Mesh ROADM	✓	**x**

TABLE 1 — Higher Order Modulation vs. Higher Baud Rates

Higher order modulation does not change the spectral width of the wavelength, so it is able to increase the spectral efficiency and is equally applicable to fixed grid and flexi-grid WDM networks. However, it comes at a significant cost in terms of reach and therefore its ability to both decrease the cost per bit and increase the spectral efficiency is limited to shorter distances.

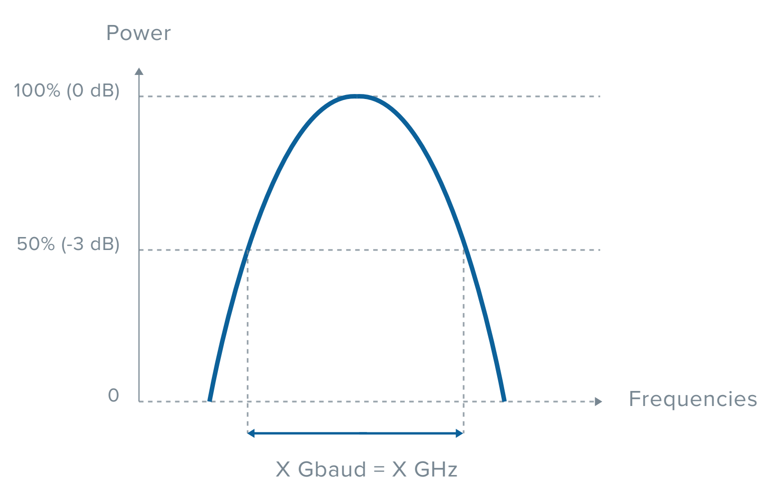

FIGURE 3 – *Symbol Rate and Channel Width*

As shown in Figure 3, the spectral width of the wavelength in GHz is equal to the symbol rate in Gbaud when measured at the 3 dB point, where the power is 50% of the peak. As the baud rate increases, the spectral width of the channel increases proportionally. Higher baud rates therefore are typically unable to increase spectral efficiency, though there can be exceptions to this rule where a higher baud rate better aligns to the available spectrum/grid structure, as we will see later in this white paper. Increasing the wavelength capacity with the baud rate however has far less impact on reach than increasing it with higher order modulation.

Higher baud rates therefore offer the best potential for reducing the cost per bit in flexi-grid DWDM networks and to a lesser extent in point-to-point fixed grid networks, though higher baud rates have a more limited role to play in 50 GHz fixed grid ROADM networks. Higher baud rates also require all the components of the optical interface including the DSP, photo detector, modulators, and A/D converters to be capable of supporting the higher bandwidth. This places a limit on the maximum baud rate that is achievable with a given set of technology and may increase the cost of the interfaces if more expensive components are required.

Cost Per Bit vs. Spectral Efficiency

Lower cost per bit and higher spectral efficiency are key requirements for evolving coherent transport. Traditionally in DWDM technology, these two moved in concert, when you decreased the cost per bit, you increased the spectral efficiency. However as the spectral width of a wavelength is equal to its baud rate, increasing the baud rate can decrease the cost per bit without increasing the spectral efficiency. In fact as a higher baud rate reduces the reach, it can even reduce the spectral efficiency. However, there are exceptions such as a point-to-point fixed grid and flexi-grid boundary case where a higher baud rate aligns better and wastes less spectrum.

Doubling Wavelength Capacity: Modulation vs. Baud Rate

To understand the relative advantages of higher order modulation and higher baud rates, let us look at what happens if we double the wavelength capacity, starting with a 100G wavelength based on QPSK modulation. The exact baud rate is also impacted by the FEC overhead. For example, 100G QPSK with 7% FEC will have a lower baud rate at around 28 Gbaud compared to 100G with 27% SD-FEC close to 34.5 Gbaud. For the purposes of this example, and in the rest of this white paper, we will use 34.5 Gbaud.

	QPSK 		16QAM 
Bits per Symbol	2	x2 →	4
Constellation Points	4	x4 →	16

TABLE 2 — QPSK vs. 16QAM: Bits per Symbol & Constellation Points

One option for 200G is to double the bits per symbol by going from QPSK (2 bits per symbol) to 16QAM (4 bits per symbol). However, as shown in Table 2, this increases the number of constellation points by a factor of four, making the wavelength approximately four times (~6 dB) more sensitive to noise and nonlinearities. The reach is therefore decreased to approximately 25% of the 100G QPSK wavelength, as shown in Figure 4. However, ignoring any grid constraints, the spectral width of the channel stays at 34.5 GHz, so we have doubled the spectral efficiency from 2.9 bits/s/Hz to 5.8 bits/s/Hz.

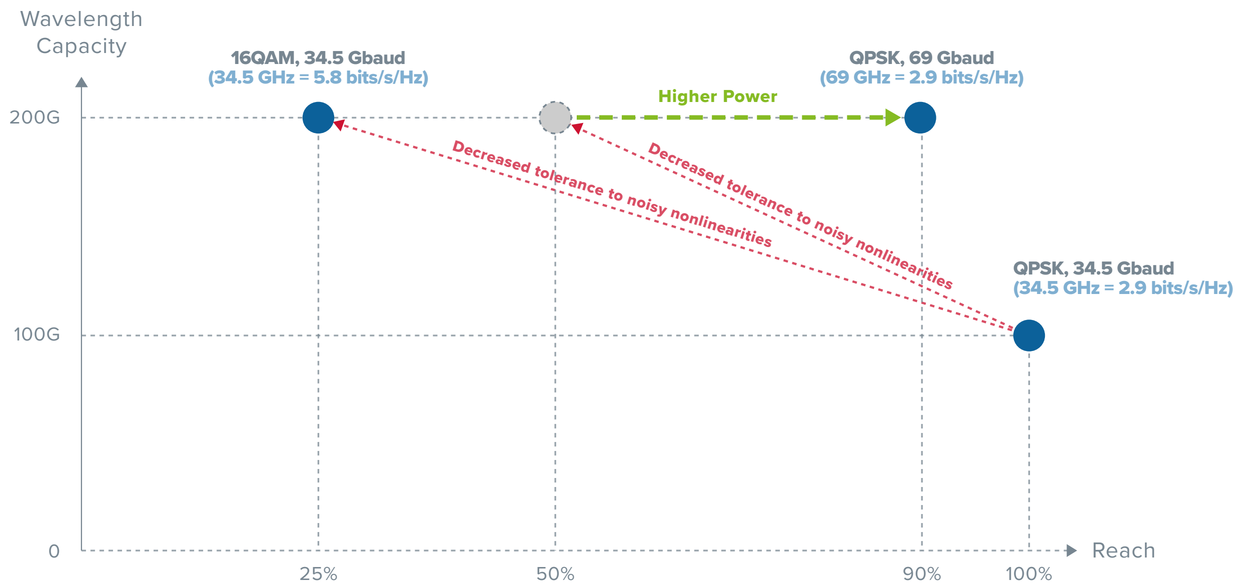

FIGURE 4 – *Impact of Doubling Wavelength Capacity*

On the other hand, if we go to 200G by doubling the baud rate to 69 Gbaud and keeping the modulation at QPSK, we double the sensitivity to noise and nonlinearities, which by itself would reduce the reach to 50%. However, we get most of this reach back with the ability to increase the total power of the wavelength. Increasing the power provides more tolerance to noise but increases nonlinear penalties such as self phase modulation and cross phase modulation. The reason we can increase the total power of the 69 Gbaud wavelength relative to the 34.5 Gbaud wavelength is that the power is now spread over a wider spectrum (69 GHz vs. 34.5 GHz) resulting in lower power density, as shown in Figure 5.

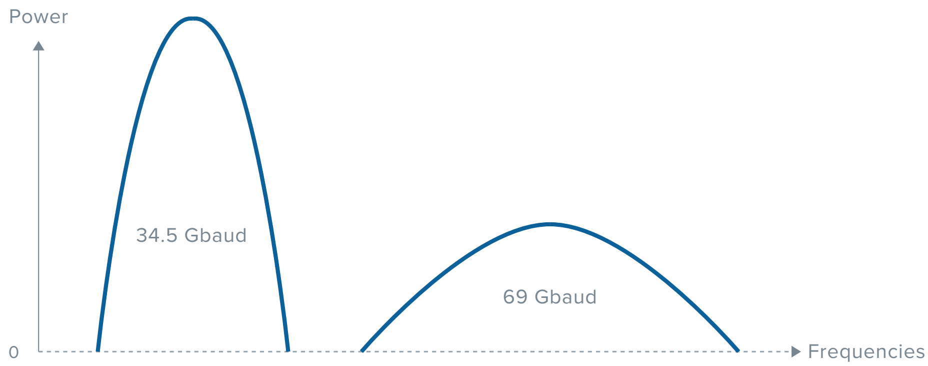

FIGURE 5 – *34.5 Gbaud and 69 Gbaud with the Same Total Power*

This typically puts the reach of the 69 Gbaud 200G at around 90% of the original 34.5 Gbaud 100G wavelength, with 3.6x the reach of the 200G based on 34.5 Gbaud and 16QAM. The price to be paid for this increase in reach is a loss of spectral efficiency relative to the higher order modulation option. The 69 Gbaud 200G wavelength now has the same spectral efficiency as the 100G 34.5 Gbaud wavelength (2.9 bits/s/Hz) and half that of the 16QAM 200G wavelength (5.8 bits/s/Hz).

HIGHER BAUD RATES AND FLEXI-GRID NETWORKS

In a flexi-grid network, the best strategy for the majority of cases will be to maximize the baud rate, then run at the highest modulation the reach requirement will allow. This will deliver the lowest cost per bit and, in the majority of cases, the best spectral efficiency, as the higher baud rate does not reduce the spectral efficiency. There will however be boundary cases where a lower baud rate can deliver better spectral efficiency. One of these boundary cases is reach, the other is the size of the block of available spectrum.

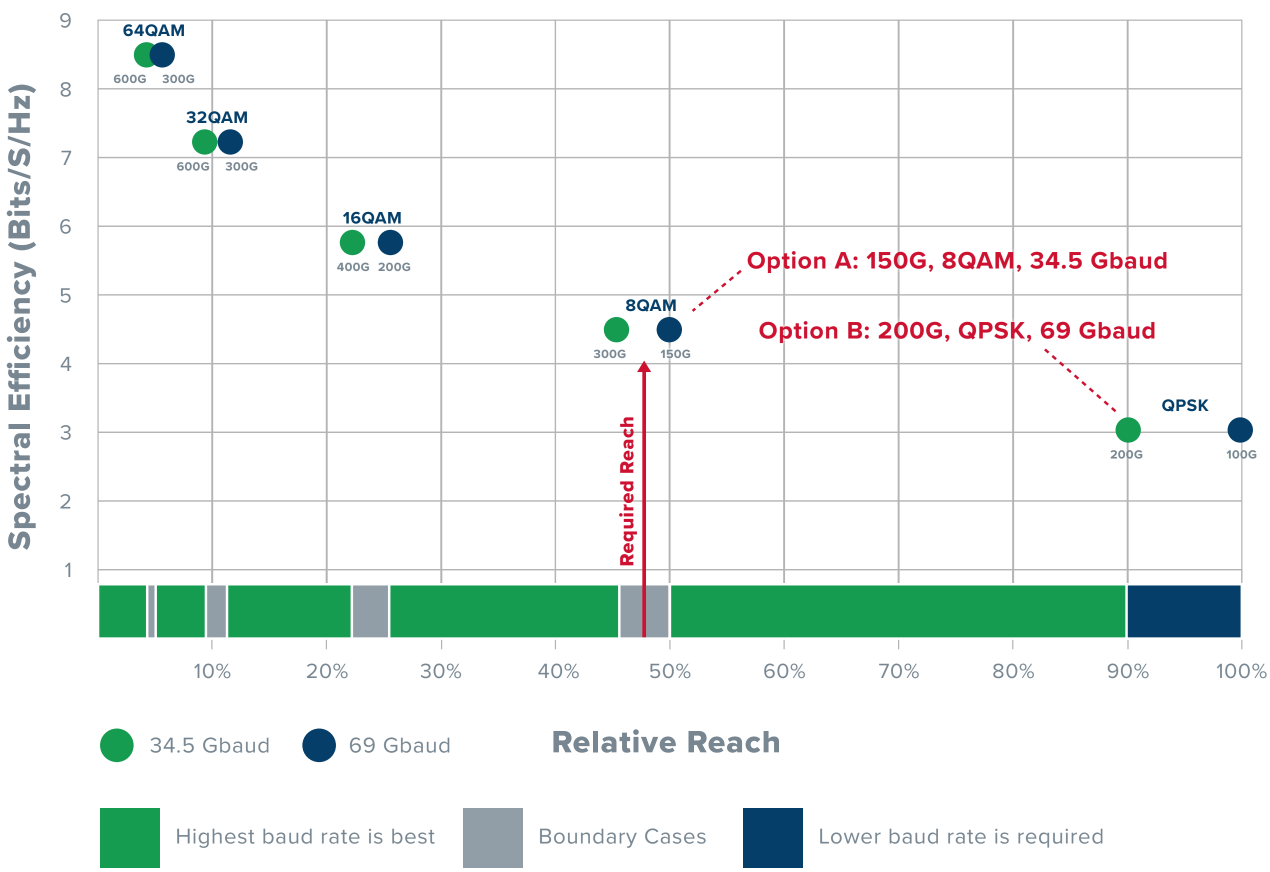

FIGURE 6 – *Spectral Efficiency vs. Reach*

Flexi-Grid Boundary Cases: Reach

Figure 6 shows the spectral efficiency versus the reach for 34.5 Gbaud and 69 Gbaud, and for QPSK, 8QAM, 16QAM, 32QAM, and 64QAM. As the figure shows, for a given modulation the baud rate does not change the spectral efficiency, however, doubling the baud rate results in a reach reduction of around 10%. For the majority of reach requirements, the higher baud rate will therefore result in the lowest cost per bit and the same spectral efficiency as the lower baud rate. For each modulation, there will be boundary cases where the higher baud rate cannot meet the reach requirement while a lower baud rate or lower modulation could.

For example if the reach requirement is just above the maximum that can be supported with 69 Gbaud 8QAM, as shown in Figure 6, then we have two options. Option A is to decrease the baud rate to 34.5 Gbaud, maintaining a spectral efficiency of 4.3 bits/s/Hz but increasing the cost per bit with only 150G per interface. Option B would be to keep the baud rate at 69 Gbaud but drop the modulation to QPSK, resulting in lower spectral efficiency of 2.9 bits/s/Hz but with 200G per interface delivering a better cost per bit than option A. As we will see later in the white paper, baud rate flexibility and hybrid modulation are key to providing a better set of options for such boundary cases.

Flexi-Grid Boundary Cases: Spectrum Block Size

A second boundary case is size of the available block of spectrum for a super-channel. For example if we have a 200 GHz block of spectrum to fill along a particular path between two points in the network and the reach requirement is just below the 8QAM range, we would have four options as shown in Figure 7. In this example, four channels of 200G 8QAM at 46 Gbaud gives the best spectral efficiency at 4 bits/s/Hz, while two channels at 69 Gbaud still gives the lowest cost per bit with 300G per interface. However, three channels of 250G at 57.5 Gbaud may give the best trade-off between cost per bit and spectral efficiency.

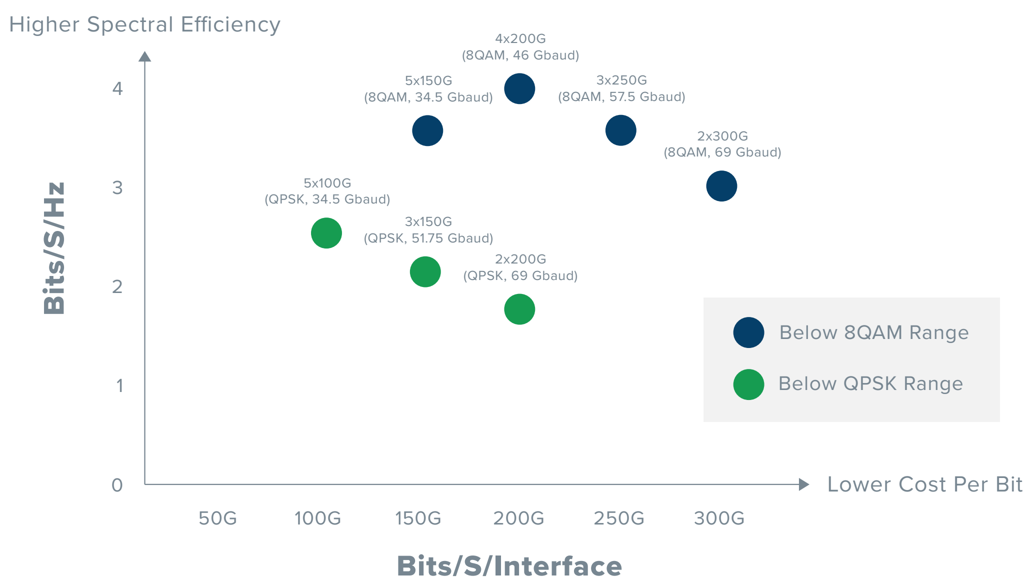

FIGURE 7 – *200 GHz Boundary Examples*

As another example, based on the options for 200 GHz at just below the QPSK reach range, also shown in Figure 7, we can achieve the maximum spectral efficiency with five channels of 100G at 34.5 Gbaud. The lowest cost per bit can be achieved with two channels of 200G at 69 Gbaud, while three channels of 150G at 51.75 Gbaud provides an intermediate option in terms of both cost per bit and spectral efficiency.

50 GHz FIXED GRID ROADM NETWORKS

Higher baud rates have limited applicability to 50 GHz fixed grid ROADM networks for a couple of reasons. First, as discussed previously, increasing the baud rate results in a proportional increase in the spectral width of the channel. Second, legacy Wavelength Selective Switches (WSSs) have a limited passband for each channel and this passband reduces significantly as the number of WSSs in the path increases, as shown in Figure 8.

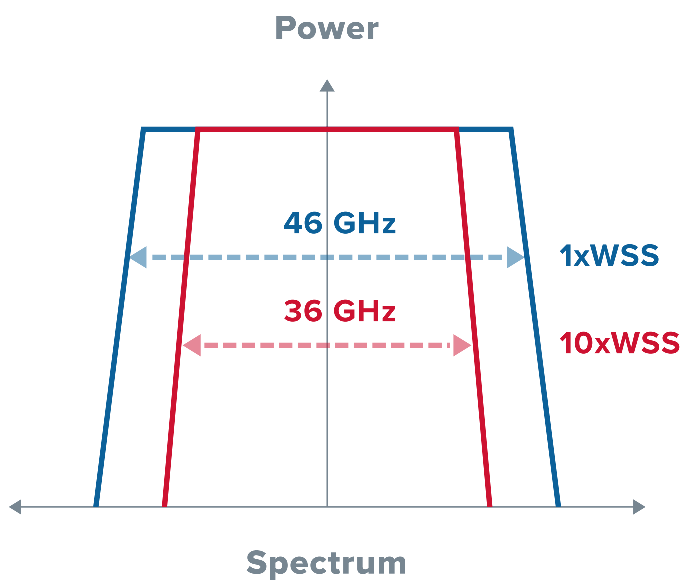

FIGURE 8 – *Filter Narrowing and WSS Cascades*

A typical legacy 50 GHz single WSS has a passband of around 46 GHz per channel. However, any differences in width, center wavelength, or shape of the passband will also cause the effective width to reduce in an effect called filter narrowing. After only ten WSSs (e.g., a colored/directional add, four route and select ROADMs, and a colored/directional drop), the effective width would be reduced by around 10 GHz to around 36 GHz, limiting the baud rate to under 35 Gbaud. In 50 GHz mesh ROADM networks, the only real role for higher baud rates is increasing the FEC.

50 GHz FIXED GRID POINT-TO-POINT NETWORKS

50 GHz mux/demux filters based on a 50 GHz interleaver typically have a passband of around 46 GHz per channel, similar to a legacy WSS. Baud rates of up to 46 Gbaud can therefore be used for point-to-point 50 GHz fixed grid applications. This can enable a lower cost per bit and single channel 400G wavelengths based on 46 Gbaud and 64QAM. In this scenario, a higher baud rate does actually increase the spectral efficiency, as less spectrum within the 50 GHz grid is wasted, as shown in Figure 9.

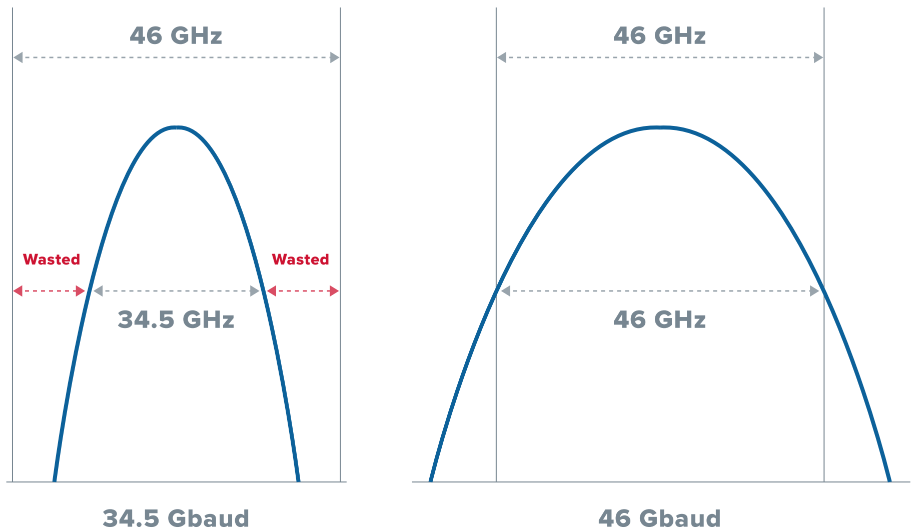

FIGURE 9 – *Higher Baud Rates and Increased Spectral Efficiency in Point-to-point Fixed Grid*

INTRODUCING INFINERA CLOUDWAVE T OPTICS

Coherent technology has undergone a number of evolutions, as shown in Table 3. First, it went from 40G (Generation 1) to 100G (Generation 2) and then from hard decision FEC (Generation 2) to soft decision FEC (Generation 3). Then it evolved from single rate interfaces (Generation 3) to flexi-rate interfaces (Generation 4) with a baud rate still in the 30 Gbaud range but with support for additional modulation schemes enabling support for multiple line rates, typically 100G (QPSK), 150G (8QAM), and 200G (16QAM).

	Gen 1	Gen 2	Gen 3	Gen 4	Gen 5
Approximate Timeframe	**2008-2009**	**2010-2011**	**2012-2014**	**2015-2017**	**2017-2019**
Wavelength Speed	40G	100G	100G	100G/150G/200G	100G to 600G
Baud Rate	12.5 Gbaud →	~30 Gbaud	30 Gbaud	~30 Gbaud →	+Higher Baud Rates
Forward Error Correction	Hard Decision	Hard Decision →	Soft Decision	Soft Decision	Soft Decision
Modulation Formats	QPSK	QPSK	QPSK →	QPSK/8QAM/16QAM →	+32QAM, 64QAM

TABLE 3 — Coherent Technology Evolution

Now a fifth generation of coherent technology is emerging with support for higher baud rates and additional modulation formats. Infinera CloudWave T Optics, Infinera's fifth generation coherent optics technology, leverages DSP technology based on a 16 nm CMOS process. This offers significant advantages over fourth and fifth generation coherent based on 28 nm technology. The step change increase in transistor density of 16 nm enables the DSP to do far more while consuming far lower power and minimizing footprint. 16 nm is a key enabler for 69 Gbaud and other capabilities including linear and nonlinear compensation and enhanced FEC. Key features of Infinera CloudWave T Optics are listed in the following paragraphs.

69 Gbaud and 600G Wavelengths

Infinera CloudWave T Optics technology delivers an industry-leading baud rate of 69 Gbaud enabling lower costs per bit to be achieved in flexi-grid networks. With 64QAM modulation, 69 Gbaud enables single channel 600G wavelengths delivering the lowest possible cost per bit at distances of up to 300 km. 69 Gbaud also enables 200G at distances of up to 4,000 km in terrestrial networks.

Flexible Baud Rate Setting

A key differentiator for Infinera CloudWave T is its ability to support a wide range of different baud rates between 28 Gbaud and 69 Gbaud. As shown in Figure 10, baud rates at the bottom of this range are needed for 50 GHz fixed grid ROADM networks (28~35 Gbaud) for the maximum possible unregenerated reach (34~35 Gbaud) and for interoperability with third parties (28 Gbaud with 7% staircase FEC). Baud rates up to 46 Gbaud also have a role to play in point-to-point 50 GHz fixed grid networks by lowering the cost per bit and increasing spectral efficiency including support for single channel 400G.

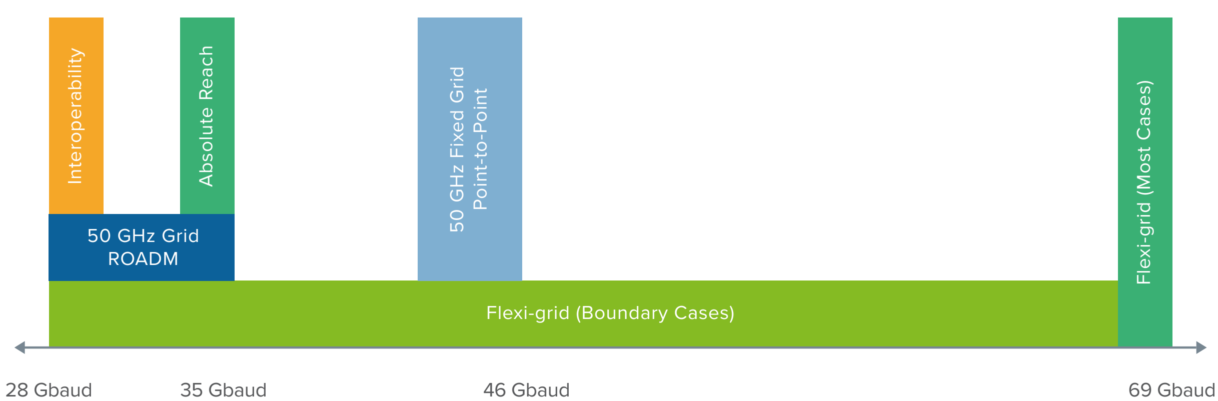

FIGURE 10 – *Baud Rate Use Cases*

In the majority of flexi-grid cases, 69 Gbaud will give the lowest cost per bit and the best spectral efficiency. However as discussed previously, there will be boundary cases where the ability to select intermediate baud rates will provide options for the best spectral efficiency or the best trade-off between cost per bit and spectral efficiency. Examples of potentially useful baud rates that align to 50G capacity increments based on standard modulation are shown in Table 4.

	QPSK	8QAM	16QAM	32QAM	64QAM
100G	34.5 Gbaud				
150G	51.75 Gbaud	34.5 Gbaud			
200G	69 Gbaud	46 Gbaud	34.5 Gbaud		
250G		57.5 Gbaud	43.125 Gbaud	34.5 Gbaud	
300G		69 Gbaud	51.75 Gbaud	41.4 Gbaud	34.5 Gbaud
350G			60.375 Gbaud	48.3 Gbaud	40.25 Gbaud
400G			69 Gbaud	55.2 Gbaud	46 Gbaud
450G				62.1 Gbaud	51.75 Gbaud
500G				69 Gbaud	57.5 Gbaud
550G					63.25 Gbaud
600G					69 Gbaud

TABLE 4 — Intermediate Baud Rate Examples

Additional Modulation Formats

The option of additional higher modulation formats also has a role to play in terms of increasing spectral efficiency and lowering the cost per bit at shorter distances. Infinera CloudWave T adds 32QAM and 64QAM to the QPSK, 8QAM, and 16QAM available with the previous generation of Infinera CloudWave Optics. Furthermore, Infinera CloudWave T is able to support time domain hybrid modulation. This hybrid modulation mixes symbols with different modulation on the same wavelength. For example, hybrid 8QAM/QPSK could alternate 8QAM symbols and QPSK symbols delivering a modulation with a capacity and spectral efficiency that is the average of the two individual modulations. As the input signal quality to the FEC block is the average of the lower and higher order modulation, the reach will get near to the average of the two individual modulations. In addition, the ratio of symbols can change (for example, 2xQPSK symbols then 1x8QAM) to provide even more granularity.

For example, as shown in Figure 11, hybrid modulation could be used to add 50 Gbps granularity to 69 Gbaud enabling the optimal cost per bit and spectral efficiency in more cases. Referring to the example in Figure 5, hybrid modulation provides a better alternative to option B (200G QPSK 69 Gbaud with 2.9 bits/s/Hz) of 250G QPSK/8QAM 69 Gbaud with 3.625 bits/s/Hz.

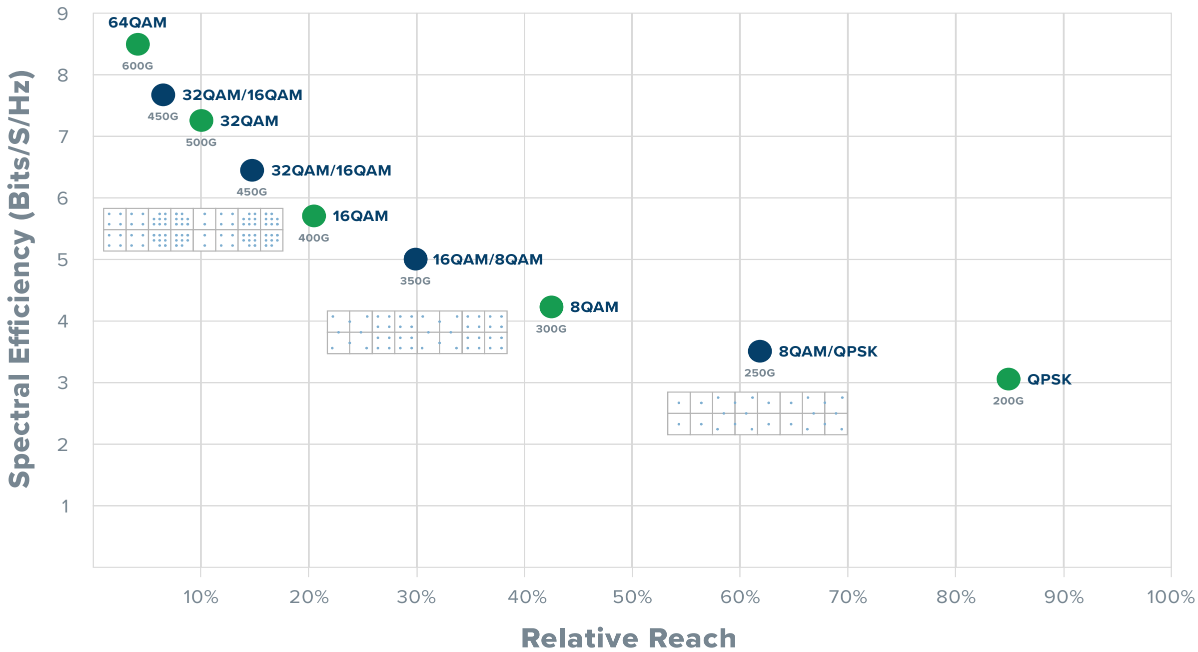

FIGURE 11 – *69 Gbaud with Hybrid Modulation*

Enhanced Forward Error Correction

FEC provides an additional lever to decrease the cost per bit and increase the spectral efficiency of the network. By reducing OSNR requirements, FEC enables better reach or more capacity with a higher baud rate and/or higher modulation. Infinera CloudWave T supports a 27% soft decision FEC, which together with other enhancements, reduces the OSNR requirement by up to 1 dB, relative to the previous generation of Infinera CloudWave Optics with 25% soft decision FEC. In addition, Infinera CloudWave T supports 15% FEC and a 7% staircase FEC for interoperability with third-party interfaces in metro applications.

INFINERA CLOUDWAVE T BENEFITS

Cost Per Bit and Spectral Efficiency

As discussed previously, 69 Gbaud, baud rate flexibility, additional higher order modulation schemes, hybrid modulation, and enhanced FEC all have the potential to lower cost per bit and/or increase spectral efficiency. However, realizing this potential requires optimally setting all the parameters including baud rate range, modulation, FEC, power levels, channel frequency, and super-channel schemes. This capability is provided by Infinera Aware Technology. As shown in Figure 12, Infinera Aware Technology comprises two key elements, the Optical Performance Engine (OPE) and the Margin Processing Engine (MPE). The MPE is able to measure the real-time residual margin of each channel while the OPE is able to generate the valid parameter options for each channel enabling the best options to be automatically selected by the NMS, SDN, or ASON/GMPLS control plane. For more details on Infinera Aware Technology, refer to the Infinera white paper *Evolving the Awareness of Optical Networks*.

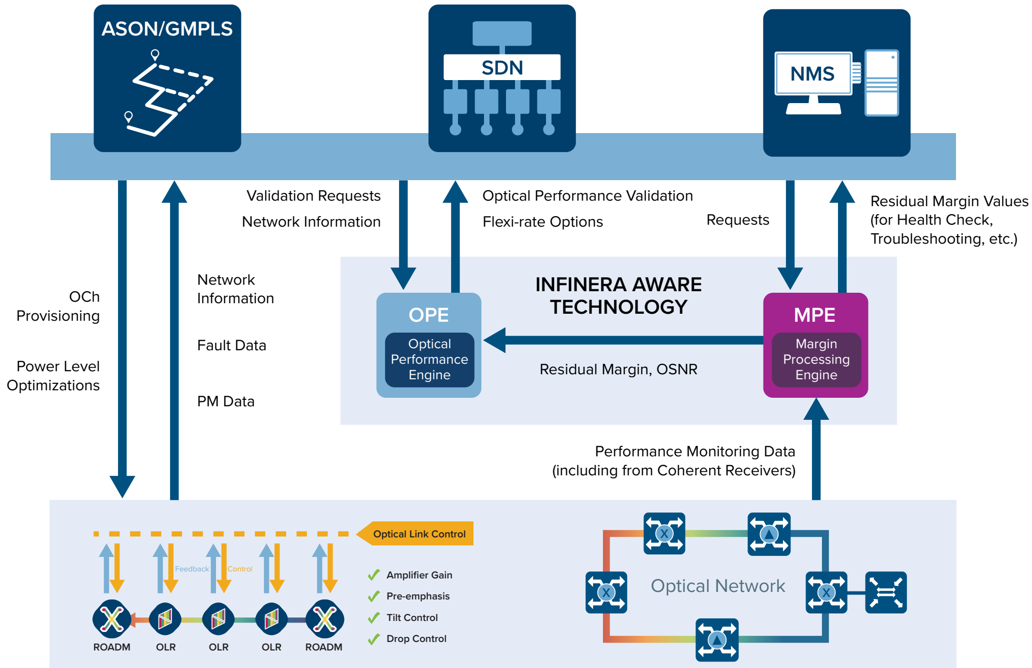

FIGURE 12 – *Infinera Aware Technology*

Minimize Power Consumption and Footprint

The higher density and lower power consumption of 16 nm technology enables Infinera CloudWave T to deliver a 3x improvement in interface capacity with the same power consumption as the previous generation of Infinera CloudWave Optics. Infinera CloudWave T reduces the W per Gbps by two thirds from 0.45 W per Gbps to 0. 15 W per Gbps, which is significantly lower than the figures between 0.4 W per Gbps and 0.74 W per Gbps currently claimed for fifth generation coherent by other equipment vendors.

Optimize Performance for Any Open Line System

As previously discussed, OLSs are a key trend in the optical networking industry. However, OLSs can have a wide variety of capabilities such as the DWDM grid, amplifier performance, power monitoring, and link control. Baud rate and modulation flexibility enables the performance of any OLS to be maximized regardless of its capabilities. Furthermore, 28 Gbaud, QPSK, and 7% staircase FEC provide an option for interoperability with third-party and pre-fifth generation coherent interfaces over shorter distances.

Support for New Client Types

The ability to support 100G to 600G per wavelength in 50G increments provides an ideal solution for supporting 200 GbE and 400 GbE clients. FlexE and FlexO enable client bandwidth to adjust to the maximum wavelength capacity enabled by the combination of Infinera CloudWave T and Infinera Aware Technology.

SUMMARY

Optical equipment vendors are evolving coherent technology in order to meet network operator demands for increased capacity at a lower cost and to support new client interfaces such as 400 GbE and FlexE. In addition to more modulation schemes, the fifth generation of coherent provides the option of higher baud rates. While for flexi-grid networks, as a general rule, higher is better, there will be boundary scenarios where lower baud rates can provide the better result in terms of spectral efficiency. And while 50 GHz point-to-point networks can support baud rates up to 46 Gbaud, 50 GHz mesh ROADM networks will be typically limited to under 35 Gbaud. Infinera CloudWave T provides the ability to select from a wide range of baud rate options between 28 Gbaud and 69 Gbaud while delivering benefits including 600G wavelengths and power consumption down to 0.15 W per Gbps. The combination of Infinera Aware Technology and Infinera CloudWave T enables cost and spectral efficiency to be optimized over a wide range of use cases.